

WHAT IS CLAIMED IS:

1. A node comprising:
  - a processor;
  - 5 a memory;
  - wherein the processor is operable to execute program instructions stored in the memory to implement:
    - receiving a request from a client application, wherein the request requires a transaction;
    - 10 in response to the request, sending a first message to a plurality of participant nodes participating in the transaction;
    - in response to receiving a reply to the first message from at least a quorum of the participant nodes, sending a second message to the plurality of participant nodes;
    - in response to receiving a reply to the second message from at least a quorum
    - 15 of the participant nodes:
      - returning success to the client application; and
      - sending a third message to the plurality of participant nodes, wherein the third message instructs the participant nodes to commit the transaction.
- 20 2. The node of claim 1,
  - wherein said sending the third message to the plurality of participant nodes comprises sending the third message after said returning success to the client application.
3. The node of claim 1,
  - 25 wherein the first message comprises a message requesting each of the participant nodes to reply by indicating whether they can commit the transaction;
  - wherein said receiving the reply to the first message from at least a quorum of the participant nodes comprises receiving a reply indicating an ability to commit the transaction from at least a quorum of the participant nodes.

30

4. The node of claim 1,  
wherein the second message comprises a message requesting each of the  
participant nodes to enter a state indicating that the transaction is to be committed;  
wherein said receiving the reply to the second message from at least a quorum  
5 of the participant nodes comprises receiving a reply indicating movement to the state  
indicating that the transaction is to be committed from at least a quorum of the participant  
nodes.

5. The node of claim 1,  
10 wherein the first message corresponds to a message for a first phase of a three-  
phase commit protocol.

6. The node of claim 1,  
wherein the second message corresponds to a message for a second phase of a  
15 three-phase commit protocol.

7. The node of claim 1,  
wherein the processor is operable to execute program instructions stored in the  
memory to further implement:  
20 forgetting about the transaction after said sending the third message.

8. The node of claim 1,  
wherein the node does not expect a reply to the third message.

9. The node of claim 1,  
25 wherein each participant node commits the transaction in response to receiving  
the third message but does not return a reply to the third message.

10. The node of claim 1,  
30 wherein the request comprises a request to update a file;

wherein the request requires a transaction to update multiple replicas of the file, wherein each replica is located on a participant node.

11. The node of claim 1,  
5 wherein the node is a node in a peer-to-peer network;  
wherein the peer-to-peer network implements a distributed file sharing system.

12. A node comprising:  
10 a processor;  
a memory;  
wherein the processor is operable to execute program instructions stored in the memory to implement:  
receiving a request from a client application, wherein the request requires a  
15 transaction;  
in response to the request, sending a first message to a plurality of participant nodes participating in the transaction;  
in response to receiving a reply to the first message from at least a quorum of the participant nodes, sending a second message to the plurality of participant nodes;  
20 in response to receiving a reply to the second message from at least a quorum of the participant nodes, sending a third message to the plurality of participant nodes, wherein the third message instructs the participant nodes to commit the transaction;  
wherein the node does not expect a reply to the third message from any participating node.

25  
13. The node of claim 12,  
wherein the node forgets about the transaction after said sending the third message.

30

14. A carrier medium comprising program instructions executable to implement:

receiving a request from a client application, wherein the request requires a transaction;

5 in response to the request, sending a first message to a plurality of participant nodes participating in the transaction;

in response to receiving a reply to the first message from at least a quorum of the participant nodes, sending a second message to the plurality of participant nodes;

10 in response to receiving a reply to the second message from at least a quorum of the participant nodes:

returning success to the client application; and

sending a third message to the plurality of participant nodes, wherein the third message instructs the participant nodes to commit the transaction.

15 15. The carrier medium of claim 14,  
wherein said sending the third message to the plurality of participant nodes comprises sending the third message after said returning success to the client application.

16. The carrier medium of claim 14, wherein the program  
20 instructions are further executable to implement:  
forgetting about the transaction after said sending the third message.

17. The node of claim 1,  
wherein the request comprises a request to update a file;  
25 wherein the request requires a transaction to update multiple replicas of the file, wherein each replica is located on a participant node.

18. A carrier medium comprising program instructions executable to  
30 implement the method of:

a first node receiving a request from a client application, wherein the request requires a transaction;

in response to the request, the first node sending a first message to a plurality of participant nodes participating in the transaction;

5 each of the plurality of participating nodes replying to the first message by indicating an ability to commit the transaction;

the first node sending a second message to the plurality of participant nodes;

each of the plurality of participating nodes replying to the second message by indicating entrance to a state indicating that the transaction is to be committed;

10 in response to receiving replies to the second message from at least a quorum of the participant nodes, the first node returning success to the client application; and

the first node sending a third message to the plurality of participant nodes, wherein the third message instructs the participant nodes to commit the transaction.

15 19. The carrier medium of claim 18,

wherein said sending the third message to the plurality of participant nodes comprises sending the third message after said returning success to the client application.

20 20. The carrier medium of claim 18, further comprising program instructions to implement:

each participant node committing the transaction in response to receiving the third message;

wherein the participant nodes do not send a reply to the third message to the first node.

25

21. A system comprising:

a network;

a plurality of computing nodes coupled via the network;

wherein the plurality of nodes store a plurality of replicas for each of a plurality of objects;

wherein the plurality of nodes are operable to perform a distributed transaction protocol to provide atomicity for an update operation for updating a plurality of replicas of a first object;

wherein said performing the distributed transaction protocol comprises a first node sending a first message, a second message, and a third message to nodes on which the plurality of replicas of the first object are located;

wherein said performing the distributed transaction protocol comprises the nodes on which the plurality of replicas of the first object are located returning a reply to the first message and the second message but not the third message.

22. The system of claim 21,

wherein said not returning a reply to the third message enables the plurality of replicas of the first object to be updated with increased efficiency compared to a three-phase commit protocol.